CS 7150: Deep Learning — Fall 2024— Paul Hand

HW 3

Due: Friday November 15, 2024 at 9:00 PM Eastern time via Gradescope

Names: [Put Your Names Here]

You will submit this homework in groups of up to 3. You may consult any and all resources. Note that some of these questions are somewhat vague by design. Part of your task is to make reasonable decisions in interpreting the questions. Your responses should convey understanding, be written with an appropriate amount of precision, and be succinct. Where possible, you should make precise statements. For questions that require coding, you may either type your results with figures into this tex file, or you may append a pdf of output of a Jupyter notebook that is organized similarly. You may use code available on the internet as a starting point.

Question 1. *Image denoising by ResNets*

Consider the following denoising problem. Let x be a color image whose values are scaled to be within [0,1]. Let y be a noisy version of x where each color channel of each pixel is subject to additive Gaussian noise with mean 0 and variance σ^2. You will need to clip the values of y in order to ensure it is a valid image. The denoising problem is to estimate x given y.

(a) Look up the definition of Peak Signal-to-Noise Ratio (PSNR). Determine what value of σ corresponds to an expected PSNR between x and y of approximately 20 dB.

Response:

(b) Create a noisy version of the CIFAR-10 training and test dataset, such that it has additive Gaussian white noise with PSNR approximately 20 dB. Show several pairs of images and their noisy version.

Response:

(c) Train a ResNet to denoise noisy CIFAR-10 images. Your net should take a noisy 32x32 px image as an input, and it should output a denoised 32x32 px image. Specify the architecture and training details of your network. Determine the mean and standard deviation of the recovery PSNRs over the noisy test set. Visually show the performance on three noisy test images.

Response:

(d) Repeat the previous task but without the skip connections in your model.

Response:

Question 2. *Adversarial examples*

Obtain a pretrained classifier for ImageNet, such as AlexNet or ResNet101 from TorchVision. Using a camera, take a picture of an object that belongs to one of the ImageNet classes. Resize it as appropriate. Select a target class that is different from the image's true class. Compute an adversarial perturbation that is barely perceptible to the human eye and that results in the image being misclassified as the target class. Clearly state the method that you used to generate the perturbation. Show the underlying image, the perturbed image, the perturbation, the classifier's confidence for the underlying image, and the classifier's confidence for the perturbed image.

Response: